

System For Knowledge Transfer In A Group Setting

BACKGROUND

Software that supports cooperative work, often called groupware, generally employs work surfaces that function according to the “What you see is what I see” (“WYSIWIS”) paradigm. This means that whenever a group of users is using such software, each has the same view. Anything written by one of the users is immediately seen by everyone else. Although the WYSIWIS paradigm has been the central idea behind most shared work surface groupware, deviations from a strict interpretation of some aspects of this paradigm have sometimes been considered.

For example, users may be distracted by a strict adherence to the WYSIWIS “simultaneity constraint.” Consider a group member who is drawing a rectangle on the shared work surface. The other group members may find it distracting to watch the fine grain updates that transpire as the rectangle is sized and placed. Therefore, it may be more appropriate for the system to delay the process of updating the other displays until the size and placement of the rectangle have been finalized, even though this implies the users will not always have identical views of the shared work surface.

Another example of an advantageous relaxing of strict adherence to the WYSIWIS paradigm is permitting the shared work surface to exceed the size of what can be displayed at one time, and permitting individual users to select which portion is displayed at their workstation.

In many situations the members of a group employing a groupware system do not all have a common role. For example, one valuable use of groupware is as a teaching aid in a classroom. In this situation, a teacher or professor has the primary responsibility for

selecting the subject matter and directing the discussion. It is therefore advantageous for this person to have additional means of reviewing and selecting the contributions from the other group members for the purpose of providing those that are likely to be most helpful to the discourse.

What is needed, therefore, is a system for sharing knowledge that uses a relaxed version of the WYSIWIS paradigm to further facilitate the efficient transfer of useful information, while filtering out information that is not useful or distracting. The present invention is directed towards meeting this need.

SUMMARY OF THE INVENTION

A first embodiment system for knowledge transfer in a group setting comprises: a plurality of participant work areas and a moderator work area. Each of the plurality of participant work areas has at least one corresponding participant input-device, and each of the participant input-devices is adapted to define participant images that are then included on the corresponding participant work area. The moderator work area comprises at least one moderator input-device, the at least one moderator input-device being adapted to define moderator images that are then included on the moderator work area and to select moderator images that are then simultaneously included on each of plurality of participant work areas. The moderator input-device is further adapted to select participant images from any of the plurality of participant work areas that are then included on the moderator work area.

A second embodiment system for knowledge transfer in a group setting comprises: a plurality of participant work stations and a moderator work station. Each of the participant work stations comprises: a participant work area; a participant input device, and a participant display element. The participant input-device is adapted to define participant images that are then included on the corresponding participant work area, the participant images being able to be defined by at least text objects, draw objects, erase objects, and bitmap objects. The participant display element is adapted to display the corresponding participant work area. The moderator work station comprises: a moderator work area, at least one moderator input-device, and a moderator display element. The at least one moderator input-device is adapted to define moderator images

that are then included on the moderator work area and on each of plurality of participant work areas, the moderator images being able to be defined by at least text objects, draw objects, erase objects, and bitmap objects. The moderator display element is adapted to display the moderator work area. The moderator input-device is further adapted to select participant images from any of the plurality of participant work areas that are then included on the moderator work area. The work areas comprise a scroll of panels, each panel being permitted to have a size exceeding that which can be displayed on the display elements. The plurality of display elements are substantially located to permit a user to simultaneously view one of the display elements and to hear substantially every other user that is viewing another of the display elements. Additional participant work stations can be added to the system while the system is in use.

A third embodiment system for knowledge transfer in a group setting comprises a plurality of work areas, each comprising a public scroll and a private scroll, and wherein one of the plurality of work areas is a moderator work area.

A fourth embodiment comprises a network of computers programmed for knowledge transfer in a group setting, the network comprising: a plurality of participant workstations and a moderator workstation. Each of the participant workstations is programmed to provide a participant work area and has at least one corresponding participant input-device. Each of the participant input-devices is adapted to create data structures defining participant images that are then included in the participant work area. The moderator workstation is programmed to provide a moderator work area and comprises at least one moderator input-device. The at least one moderator input-device is adapted to: create data structures defining moderator images that are then included in the

moderator work area; and to select moderator images that are then simultaneously included on each of plurality of participant work areas. The moderator input-device is further adapted to select participant images from any of the plurality of participant work areas that are then included on the moderator work area.

A fifth embodiment comprises a network of computers programmed for knowledge transfer in a group setting, the network comprising: a plurality of participant workstations and a moderator workstation. Each of the participant workstations is programmed to provide a participant work area and has at least one corresponding participant input-device. Each of the participant input-devices is adapted to create data structures defining participant images that are then included on the participant work area. The moderator workstation is programmed to provide a moderator work area and comprises at least one moderator input-device. The at least one moderator input-device is adapted: to create data structures defining moderator images that are then included on the moderator work area; and to select moderator images that are then simultaneously included on each of plurality of participant work areas. The moderator work area comprises a moderator public scroll and a moderator private scroll, and each participant work area comprises a participant public scroll and a participant private scroll. Each participant workstation displays images placed on the participant's public scroll by the moderator superimposed on images placed on the participant's public scroll by the participant. The moderator input-device is further adapted to select participant images from any of the plurality of participant work areas that are then included on the moderator work area.

A sixth embodiment comprises a network of computers programmed for knowledge transfer in a group setting, the network comprising: a plurality of participant workstations and a moderator workstation. Each of the participant workstations is programmed to provide a participant work area and has at least one corresponding participant input-device. Each of the participant input-devices is adapted to create data structures defining participant images that are then included on the participant work area. The moderator workstation is programmed to provide a moderator work area and comprises at least one moderator input-device. The at least one moderator input-device is adapted to: create data structures defining moderator images that are then included on the moderator work area; and to select moderator images that are then simultaneously included on each of plurality of participant work areas. The moderator work area comprises a shared work area and a moderator private work area, and each participant work area comprises a participant public work area and a participant private work area. Each participant work station displays images placed on the shared work area superimposed on images placed on the participant public work area. The moderator input-device is further adapted to select participant images from any of the plurality of participant work areas that are then included on the moderator work area. The network can be used in a group mode and a standalone mode. The workstations are located such that a first user positioned to use a workstation and a second user positioned to use a different workstation can hear each other speak.

A seventh embodiment comprises a network of computers programmed for knowledge transfer in a group setting, the network comprising: a plurality of participant workstations and a moderator workstation. Each of the participant workstations is

programmed to provide a participant work area and has at least one corresponding participant input-device. Each of the participant input-devices is adapted to create data structures defining participant images that are then included on the participant work area. The moderator workstation is programmed to provide a moderator work area and comprises at least one moderator input-device. The at least one moderator input-device is adapted to: create data structures defining moderator images that are then included on the moderator work area; and to select moderator images that are then simultaneously included on each of plurality of participant work areas. The moderator work area comprises a shared work area and a moderator private work area, and each participant work area comprises a participant public work area and a participant private work area. Each participant workstation displays images placed on the shared work area by the moderator superimposed on images placed in the participant public work area. The moderator input-device is further adapted to select participant images from any of the plurality of participant work areas that are then included on the moderator work area. The network can be used in a group mode and a standalone mode. Every user positioned to use a workstation can hear every other user positioned to use any other workstation. Data structures defining the images are organized in notebook files comprising at least one panel and are stored as at least one object in a single panel.

An eighth embodiment comprises a network of computers programmed for knowledge transfer in a group setting, the network comprising: a plurality of participant workstations and a moderator workstation, and collision-correction functionality. Each of the participant workstations is programmed to provide a participant work area and has at least one corresponding participant input-device. Each of the participant input-devices

is adapted to create data structures defining participant images that are then included on the participant work area. The moderator workstation is programmed to provide a moderator work area and comprises at least one moderator input-device, the at least one moderator input-device being adapted to: create data structures defining moderator images that are then included on the moderator work area, and to select moderator images that are then simultaneously included on each of the plurality of participant work areas.

The moderator work area comprises a shared work area and a moderator private work area, and each participant work area comprises a participant public work area and a participant private work area. Each participant workstation displays images placed on the shared work area by the moderator superimposed on images placed on the participant public work area. The participant work area comprises a participant public scroll and a participant private scroll, the participant public scroll comprising data structures defining images placed on the shared work area and the participant's public work area. The moderator input-device is further adapted to select participant images from any of the plurality of participant work areas that are then included on the moderator work area. A participant must decide to permit an image to be selected by the moderator input-device before it can be included on the moderator work area. The network can be used in a group mode and a standalone mode. Every user positioned to use a workstation can hear every other user positioned to use any other workstation.

A ninth embodiment system for knowledge transfer in a group setting comprises: a plurality of participant work stations and a moderator work station. Each of the participant work stations comprises: at least one participant display device; at least one input device; a participant public scroll; and a participant private scroll. The at least one

participant input device is adapted to permit the participant to create data structures on the participant public scroll and participant private scroll defining images that are displayed on the at least one participant display device. The moderator work station comprises: at least one moderator display device; at least one moderator input-device; a moderator public scroll. The at least one moderator input-device is adapted to create data structures on the moderator public scroll and each of the participant public scrolls defining images that are displayed on the at least one moderator display device and on each of the at least one participant display devices.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is an example of an image of a binary search tree that has been placed on a shared work area.

Figure 2 is an example of the image of Figure 1, superimposed with images placed on a participant public work area, including a footnote image having a corresponding text image on the participant's private work area.

Figure 3 is a diagram of certain elements of a group work area according to the present invention.

Figures 4A and 4B are diagrams showing certain dynamic sharing methods of a group work area according to the present invention.

Figure 5 is a diagram showing certain dynamic sharing methods of a group work area according to the present invention.

Figure 6 is a diagram of certain elements of a preferred embodiment user work area according to the present invention.

Figure 7 is an illustration of the evolution of an image created by adding image objects to a panel of the preferred embodiment.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

For the purposes of promoting an understanding of the principles of the invention, reference will now be made to the embodiments illustrated in the drawings and specific language will be used to describe the same. It will nevertheless be understood that no limitation of the scope of the invention is thereby intended, such alterations and further modifications in the illustrated device, and such further applications of the principles of the invention as illustrated therein being contemplated as would normally occur to one skilled in the art to which the invention relates.

A preferred embodiment system of knowledge transfer in a group setting according to the present invention employs a What-You-See-Is-What-I-See-One-Way (“WYSIWISOW”) paradigm. The WYSIWISOW paradigm is characterized by inclusion of a moderator workstation and participant workstations in a shared work surface system. The moderator workstation makes available to a user, called the moderator, additional functionality that is not available through the participant workstations to their users, known as participants.

The preferred embodiment system comprises a group work area. Figure 3 illustrates the relationship between certain elements of the group work area of the presently preferred embodiment, shown generally at 300. The group work area 300 comprises at least one moderator work area 310 and at least one participant work area 320, and a shared work area 330 and a plurality of private work areas 340, one corresponding to each workstation—including one moderator private work area 340a. The moderator adds images to the group work area 300 by adding them to the moderator

work area 310, while participants add images to the group work area 300 by adding them to their participant work area 320.

In the preferred embodiment, the moderator work area 310 comprises a shared work area 330 and a moderator private work area 340a. Similarly, in the preferred embodiment, each participant work area 320 comprises a participant public work area 345 and a participant private work area 340. When the moderator adds images to the shared work area 330, they are simultaneously displayed at each participant's workstation (and the moderator work station). When the moderator places images on the moderator private work area 340a, they are displayed at the moderator workstation, but they are not generally displayed at any participant's workstation. In certain alternative embodiments, the moderator work area 310 does not include a moderator private work area 340a.

When a participant adds images to the group work area 300, they are added either to that participant's public work area 345 or that participant's private work area 340. Images placed on a participant's public work area 345 are typically displayed superimposed on images on the shared work area 330, while images placed on a participant's private work area are generally not. (In some situations it may make a difference what order such objects are "stacked" in—e.g., when the images involve different colors. "Superimposed," in this context, is not meant to indicate which image is placed "on top" of which. In certain embodiments, the images on the participant public work area 345 are always placed "on top" of those on the shared work area 330, while in others the opposite is true. In certain other embodiments, the images on the shared work area 330 and participant public work area 345 are always "stacked" in the order in which they are created. In still other embodiments, the order in which these images are

“stacked,” can vary, for example at the selection of the participant.) The images on each participant work area 320 are typically initially displayed only at that participant’s workstation. In certain embodiments, the moderator may, however, view any of the participant work areas 320, including private work areas 340, at the moderator workstation, and may select images from a participant work area 320 to be included in the shared work area 330, thereby causing them to be displayed at each workstation.

A preferred embodiment system according to the present invention can be used in group mode or standalone mode. Typically, the system is used in group mode to initially aid in the transfer of information between users. Standalone mode can be used during follow up activities in which, for example, the moderator may check work that was produced while the system was used in group mode. During group use, images are placed on the group work area 300. The group work area 300 can be reviewed by individual users later, in standalone mode, in order to aid in retention and understanding of the information transferred. Standalone mode can also be used prior to a group use session, for example to prepare images for placement on the public scroll.

When used in group mode, the group typically discusses a subject, in order for the moderator to provide information to the participants about that subject. The moderator typically augments verbal statements by placing images on the shared work area 330 which he or she expects to aid the participants in more fully understanding those statements. For example, if the subject being discussed is chemistry, the images might include chemical symbols or equations. For another example, if the subject being discussed is the Pythagorean Theorem, the moderator might provide a diagram of a right triangle. The participants can supplement the images provided by the moderator with

images that the participant believes might be helpful for remembering the relevance of the images selected by the moderator. For example, a participant might add a caption labeling one side of the right triangle “hypotenuse.” This caption would become part of that participant’s surface 320, and would generally be displayed only at that participant’s workstation. Note that the group work area 300 can be used in environments where the workstations are located such that the users can speak directly to one another—for example, in a classroom—or in environments where some means of transmitting audio are necessary—such as a virtual classroom, in which various users are located remotely, and are connected via the Internet. The group work area 300 can also be used in environments in which the users cannot hear each other at all, such as when, for example, the users rely completely on the group work area 300 for the transmission of information during the group mode session.

It will be appreciated by those skilled in the art that, because the moderator does not normally see images placed on a participants’ surface 320, it is possible that the moderator will subsequently place an image on the shared work area 330 which is superimposed on one or more images that participants have placed on their participant public work area 345. This problem is known as “collision.” There are two general strategies for coping with collision: collision correction and collision avoidance.

Collision correction functions to permit users to work around the collisions that occur. For example, in certain alternative embodiments each workstation includes functionality that permits toggling between three view modes: one in which only the shared work area 330 is displayed, one in which only the participant public work area 345 is shown, and one in which both are superimposed. In these embodiments, although the

last of these modes might include collisions, by toggling between the other two, a participant can view any image on either their public work area 345 or the shared work area 330 without the collisions.

Another collision correction strategy involves including functionality at each workstation that permits the participants to relocate images on the participant work area 320 that have become involved in collisions. In certain embodiments, the colliding images are relocated automatically when a collision occurs.

Collision avoidance, on the other hand, functions to reduce or eliminate the probability of collisions occurring in the first place. For example, in certain embodiments a participant's public work area 345 includes an area (such as a margin or window) that does not have a corresponding location on the shared work area 330. Thus, a participant can add images to this area and be assured that no collisions will occur.

In the presently preferred embodiment, participant workstations include functionality that permits the participant to place footnote images on the public scroll 345 and private scroll 340. These footnote images provide a link between images on a participant's public work area 345 or the shared work area 330, and the private scroll 340, and vice versa. Preferably, these footnote images can also provide a link between images on a participant's public work area 345 or the shared work area 330 and some other virtual location. For example, such a footnote image might provide a link to a computer file, or to a URL for a website containing information the participant believes relevant. Preferably, participants can also add images to their private scroll 340 without regard to the danger of future collisions. Thus, a participant can, for example, place smaller annotations directly on their public work area 345, while placing larger

annotations, with their correspondingly larger danger of generating future collisions, in a footnote. In the presently preferred embodiment, the footnote images include means to display the corresponding images. For example, if the corresponding images are located in the participant's private work area 340, the footnote image preferably includes functionality to cause the participant's workstation to display the portion of the private work area 340 containing those images.

Thus, the presently preferred embodiment group work area 300 could be used to aid in teaching the participants, for example, about binary search trees. The moderator might select a place on the shared work area 330, and draw an image of a binary search tree, as shown in Figure 5, thereby placing the image shown in Figure 1 on the shared work area 330, and causing it to be displayed on each of the participant workstations. The moderator might state orally that the part of the tree containing the circled "3" is called the root, as a part of a lengthier oral description of the properties of binary search trees generally. A participant might place an annotation saying "root" in his or her participant public work area 345, to identify that portion of the binary search tree. The participant might also type a lengthier annotation on their private work area 340, with footnote symbols saying "To Private" and "To Public" and comprising hypertext links such that "clicking" on one of them (either with the drawing pad, a mouse, or some other appropriate input device) in one window causes the other window to center on the other footnote symbol. Figure 2 illustrates what this might look like on the display device of the participant's workstation. Those skilled in the art of computer operation will recognize the "scroll bars" that provide graphical information indicating how to scroll

these scrolls up and down. Note that either window could be independently scrolled to review images above or below the ones shown here.

In the presently preferred embodiment, each of the workstations is a complete and independent computer workstation sharing a network connection to each of the other workstations. Thus, in the presently preferred embodiment, each workstation includes its own processor, monitor, storage media, memory, etc. In the presently preferred embodiment, the workstations transfer information via a dynamic TCP/IP protocol, so that no shared storage is needed. Those skilled in the art will appreciate that the system need not use TCP/IP, and that any suitable protocol may be used to transmit information from workstation to workstation. In certain alternative embodiments, in order to reduce the cost of the workstations, some or all of them may include fewer than all of these features. For example, in certain embodiments, the participant workstations include a processor and monitor, but no memory or storage media. These assets are preferably located elsewhere in the network, such that the workstations can use them.

As shown in Figure 6, preferably, the moderator work area 310 and each of the participant work areas 320 are each implemented as at least one scroll, each scroll comprising a set of practically infinitely scrollable panels. In the presently preferred embodiment, for example, the moderator work area 310 comprises a moderator public scroll 610 and a moderator private scroll 612, while each participant work area 320 comprises a participant public scroll 620 and a participant private scroll 622. (In those embodiments lacking a moderator private work area 340a, the moderator private scroll 612 is excluded.) Preferably, each of the scrolls is made up of one or more panels 650 in series, similar to a slide-show. However, unlike a slide show, preferably, portions of

more than one panel 650 can be displayed at a time, so that an image can be placed straddling two panels 650 and still have each portion displayed simultaneously. “Scrollable,” in this context, means that the panels 650 can be larger than what can be displayed on the display device, but that the portion that is being displayed can be “scrolled” by means of a scroll bar or other appropriate mechanism, such as would occur to a person of ordinary skill in the art of computer use. “Practically infinite,” in this context, means limited by the ability of volatile memory to contain the entire scroll. In certain embodiments, the panels 650 can scroll vertically or horizontally. In certain embodiments, the panels can scroll both vertically and horizontally. In certain alternative embodiments, the panel’s length is limited by persistent memory, such that portions of the panel are swapped between volatile and persistent memory as necessary when the size of the panel exceeds what can be stored in volatile memory alone. Preferably, the panels have a variable size. Thus, in order to conserve space, a panel preferably is created with some relatively small size. If additional space is needed in order to accommodate larger images, the panel’s size can be increased. Thus, in the preferred embodiment, the user defines the actual size of each panel. Preferably, such additional panel size can be added even after additional panels have been added to the scroll.

In the presently preferred embodiment, each panel 650 comprises a linked-list of objects that have been added to it. Each workstation can “play back” the evolution of the panel by adding these objects, one at a time, to the display, in the order they were originally added. It will be appreciated that this means that each object generally occupies only a single panel 650. Thus, while an image can straddle two panels 650, it will generally be comprised of objects that do not. It will also be appreciated that other

data structures can also be used to store information corresponding to images in panels 650. For example, the workstations need not use an object-oriented database structure. Panels 650 could be stored in a relational database. For another example, the workstations need not even store the objects in an internal data structure. In the presently preferred embodiment, the footnote objects comprise hyperlinks that cause the corresponding annotations to be displayed when activated, for example by clicking on them with a mouse.

In the presently preferred embodiment, the moderator adds an image to the shared work area 330 by adding an object corresponding to it to the moderator public scroll 610. Information is then sent to each participant workstation causing that object to be added to each of the participant's public scrolls 620. The moderator places images on the moderator private work area 340a by adding an object corresponding to it to the moderator private scroll 612. Likewise, a participant adds images to the participant's public work area 345 or the participant's private work area 340 by adding an object corresponding to it to the participant's public scroll 620 or the participant's private scroll 622, respectively. Thus, in the preferred embodiment, images on the participant's public work area 345 are "stacked" in chronological order, without regard to who placed them there.

In the presently preferred embodiment, the moderator workstation includes an electronic whiteboard upon which the moderator can draw directly, for example, using a finger, to generate objects that are placed on the moderator's public scroll 610 or private scroll 622. In those embodiments in which at least some of the participant workstations are proximally located, the whiteboard can be oriented to permit participants to observe

while images are drawn on it. In this way, participants can receive information from non-verbal cues, such as gestures the moderator makes in relation to the images on the whiteboard. Preferably, the whiteboard can also be oriented so that the participants cannot see it, so that the moderator can prepare a complete image that can be presented to the participants as a finished product. In certain embodiments, the participant workstations are located remotely, so that participants could not see a whiteboard. In such embodiments, the whiteboard can be omitted.

In the presently preferred embodiment, the participant workstations include a pen-based video tablet, which permits the participants to place objects on the participant's public scroll 620 or private scroll 622. Preferably, the moderator workstation also includes a pen-based video tablet. Preferably, each workstation also includes a keyboard for entering text objects. Other input devices, such as mice, optical scanners, etc., can also be advantageously included at one or more of the workstations. In those embodiments in which the moderator work area 310 includes a moderator private scroll 612, the moderator workstation advantageously includes at least two monitors—one for displaying the private work area 340a, and one for displaying the shared work area 330.

The combination of a display device and a drawing input device (such as a whiteboard or pen-based video tablet) is sometimes referred to as a drawing surface. The user can create an image on the group work area 300 by drawing on the input device, which then creates a data structure defining the image. The data structure is placed on the appropriate scroll, and the image is displayed on the display device. Of course, other types of images, including graphic images stored in a graphics file, can be placed on the group work area 300, displayed on the display devices, and combined with drawn images.

Thus, for example, a digital photograph could be displayed, and some relevant feature highlighted by drawing a circle or an arrow on the drawing surface.

A preferred embodiment system according to the present invention can use dynamic or static sharing methods to transfer information between workstations. As discussed above, the presently preferred embodiment system uses TCP/IP to permit dynamic sharing methods. Figures 4A, 4B, and 5 illustrate three such dynamic sharing methods. For example, as shown in Figure 4B, the moderator can select a panel 650 from either the moderator's public scroll 610 or the moderator's private scroll 612 to be sent to a participant's workstation. For another example, as shown in Figure 4A, a moderator might pull a panel 650 from a participant's public scroll 620 or private scroll 622 by selecting the panel 650 to pull. A message is sent to the participant's workstation, which then sends the panel 650 back to the moderator workstation without participant interaction. The returned panel 650 can be placed on the moderator's public scroll 610 or the moderator's private scroll 612, at the moderator's discretion.

Figure 5 illustrates a kind of dynamic sharing in which the participant pushes a panel 650 either from the participant's public scroll 620 or private scroll 622, which is then placed on the moderator's public scroll 610 or private scroll 612, at the moderator's discretion.

Using these methods, a moderator might place a question on a panel 650 of the moderator's public scroll 610, and request that a participant pass a panel 650 with the answer from their private scroll 620 to the moderator workstation. Alternatively, the moderator might actively select a participant's answered panel 650 and pull the panel 650 without participant interaction. Or, the moderator could pose a question on the

moderator's public scroll 610, and then "collect" the answers. The answers could be collected by copying the panels 650 directly to the moderator's public scroll 610, or by copying the panels 650 to the moderator's private scroll 612, where they can be reviewed before pasting selected answers into the moderator's public scroll 610 and sharing them with the group.

A preferred embodiment system according to the present invention can also use static sharing methods. For example, one or more scroll panels 650 can be exported and written to a panel file. A user can export panels 650 from either their public scroll 610 or 620, or their private scroll 340, to a panel file. The resulting electronic files can be copied, deleted, or imported into another session. When imported into another session, the panel file can be placed by a participant on their private scroll 622 or their public scroll 620, or by the moderator either on the moderators public scroll 610 (and, consequently, each participant's public scroll 620) or the moderators private scroll 612. Panel files can also be transferred to other users, for use in the same or other sessions.

Thus, a user can save information, in the form of files, corresponding to images on their work area 310 or 320. The resulting information can be imported in another session. A participant can place these images on their public work area 345 or their private work area 340, while the moderator can place these images on the shared work area 330 or the moderator private work area 340a.

In the presently preferred embodiment, entire sessions, or portions of sessions, can be saved in a notebook file. This permits a session to be interrupted, and later resumed. The scrolls of a single user's work area form a chapter in the notebook file. Thus, for example, a participant could save only the their own work area to a notebook

file having only a single chapter. Or, for example, the moderator could save the entire group work area 300 in a notebook file having one chapter for each user that was involved in the session. Notebook files can be opened, closed, copied, or deleted. Two notebook files can be merged into a single notebook file. Notebook files can be transferred between users. The electronic size of a notebook is typically limited only by storage space on persistent memory media.

Preferably, each step in the evolution of the group work area 300 is stored in the notebook file, so that the users can review the group discussion by watching as each new image is added, or as an existing image is altered. Figure 7 illustrates the evolution of a complex image on a surface, as it is constructed from a set of simpler images. The simplest images that can be placed on the group work area 300 are called objects. In the presently preferred embodiment, the available objects include free-form draw objects, defined shapes (such as lines, circles, rectangles, arrows, polygons, ellipses, perpendicular lines, answer boxes, etc.), typed text, and erasures. In the presently preferred embodiment, “erasure” objects actually just fill a region with the background color, so that images or portions of images are masked. They do not remove previously added objects from the data structure of the panels 650, or the group work area 300.

In the presently preferred embodiment, objects on the moderator’s public scroll 610 are conveyed to participants’ workstation in a dynamic fashion. The moderator’s objects are typically conveyed only when the objects are complete. So, for example, while the moderator is creating a draw object, the participants do not see the progress of developing the draw object on their display. After a pointer creates the draw object and

the pointer is no longer focused on the object, the object is then conveyed to all participants, and added to their public scroll 620.

The presently preferred embodiment employs the structure of panels 650 as a linked-list to capture and replay the creation of objects for both moderator and participants. The system captures information to a scroll panel 650 in a sequence of individual objects as they are applied to the scroll panel 650. The objects are part of the panel 650 and are not stored as separate files. The creation sequence of objects on a scroll panel 650 may be replayed to show the sequence in an object-by-object fashion. There is no "capture" activation required by a user since the placement of objects on the panel 650 is a native function. Thus, in the preferred embodiment, there is no need for the moderator or a participant to recognize in advance that the evolution of an image may be useful, in order to activate a capture function (i.e. a "record" function). Since the objects are embedded in a panel 650 the replay feature is natively available when panel files 650, panel messages, or notebook files are communicated.

Thus, it will be appreciated that in addition to storing images corresponding to information discussed by the group on the series of panels 650 of the group work area 300 in a pattern indicative of the order in which it was discussed, such as, for example, placing them from top to bottom on the work area as they are created (or from left to right on a horizontally scrolling surface), sets of related images can also be stored on the work surface within a single panel 650, as a series of modifications to a previous image. In this way, the notebook file can store evolution in an image that took place during the session.

For example, in a discussion about architecture, the moderator might start by showing a basic floor plan that will serve as a starting point for a number of points during

the discussion. In addition to simply drawing that basic floor plan during the discussion, it might have been drawn by the moderator before the group discussion (for example, by using the system in standalone mode), or it might have been imported as a digital image (such as from a jpeg file)—for example by being downloaded from the internet—or it might have been optically scanned in by the moderator. During the discussion, the moderator could alter the floor plan in order to show the solution to some design problem. If this solution introduced some new design problem, the moderator might then alter this new floor plan. In this way, the floor plan might go through several iterations. Rather than placing a new image on a new location on the shared work surface 300, these images could be stored as alterations in the previous image. Thus, when a user later reviewed the discussion, the basic floor plan could be shown first, and then be again replaced by each iteration of that basic design that was included by the moderator during the discussion. Note that in the preferred embodiment, the later review can be done during the ongoing session, as discussed further herein, with respect to session synchronization.

In the preferred embodiment, in which the moderator work area 310 and participant work areas 320 are implemented as a public scroll and a private scroll, these alterations are stored as additional objects in series on the panel 650, including, for example, erase objects and new draw objects.

Another example of how a set of related images might be advantageously placed in the same apparent location of the group work area 300 is animation. For example, in a discussion of a rotary combustion engine, a series of images illustrating the movement of the parts, shown in a continuous closed loop, could be included, as a means of graphically

illustrating the engine cycle. Such animations could be stored as a looped series of objects, such as bit maps. In certain embodiments, such animations could be stored as an animation object.

The presently preferred embodiment includes means to synchronize a workstation with an ongoing session. For example, a participant may be late for the start of the discussion, or may be interrupted during a session. The presently preferred embodiment includes two methods for synchronization: session synchronization and scroll panel synchronization.

In the preferred embodiment, session synchronization transmits all moderator objects from the beginning of the session up to the most recent object entered by the moderator to the moderator public scroll 610. In situations where a participant was late joining a session, session synchronization allows the participant to obtain all panels from the beginning of the session. It then positions the participant's public scroll to same location as the moderator's scroll panel.

Scroll panel synchronization allows participants that start a session with the moderator to catch up while in a session. This feature is enabled by a continuous and automatic synchronization process that occurs without moderator or participant intervention. This frees the participant to not track their public scroll 620 to the moderator's public scroll 610—that is, when the participant can display a previous panel 650 or replay the link-list comprising the present panel 650. For example, the participant can also review a prior panel 650 to add local annotations which might not have occurred to them before. During this time the moderator may continue to present new panels 650 during the session. Scroll panel synchronization repositions the participant's public scroll

620 so it is at the same location as the moderator's public scroll 610. It also ensures that all panels from the moderator's public scroll 610 are available to the participant during the period the participant's public scroll 620 was not synchronized with the moderator's public scroll 610.

In the presently preferred embodiment, the moderator workstation includes functionality for an assessment mode. Certain alternative embodiment systems include this functionality elsewhere in the system. Regardless of where the assessment mode functionality is implemented, panel files 650 can include checkpoint objects. Checkpoint objects need not have any graphic representation, though in certain embodiments they do. The checkpoint object serves as a flag in the linked-list of a panel file 650. When placed in assessment mode, replay of a panel 650 either begins or ends with the selected checkpoint.

For example, in certain embodiments, checkpoint objects can be used by the moderator to mark the end of a time limit during the discussion. For example, if a teacher wished to use the assessment mode to present and review a question to the participating students, the question might be created on a panel with 25 objects (free hand draw objects, typed text objects, shapes, etc.). A participant might begin to write an answer composed of 50 objects or so. During the process of preparing this answer, the participant might realize that their approach to the answer is incorrect, and select an erase object to cover over part of their answer. The participant then could continue to add objects to complete the answer. At some point, the teacher will want to explain the correct answer. The moderator can use the assessment mode to place a checkpoint object on each participant's public scroll 620 (by placing it on the moderator's public scroll

610), and then begin to explain the correct answer. At this point, if a participant realizes their answer is again mistaken, they can again erase (i.e. cover over) the errors and add images demonstrating the correct answer. The participant can view the corrected answer normally. However, when the teacher pulls the panel 650 and replays it to review the student's performance, the replay will stop at the checkpoint, so that the teacher knows that subsequent objects were added after the teacher began to explain the correct answer.

Note that the teacher need not pull the relevant panel 650 to review an answer. For example, the teacher could open a student's notebook file in standalone mode instead.

In certain embodiments, another type of checkpoint object can be inserted by the participant, in order to mark where the teacher should begin reviewing. For example, in the example above, the student could insert such a checkpoint object after the first erase object. When the teacher pulls and reviews that panel in assessment mode, the display would begin at that point, already displaying each of the previous objects (including the erase object), and begin adding the subsequent objects to the display one at a time. In this way, the teacher need not spend time reviewing work which the student has already recognized was in error. Note that the teacher may preferably still review the earlier work, either by overriding the student's checkpoint, or by reviewing the panel outside of assessment mode.

In the presently preferred embodiment, the moderator workstation can only pull images from a participant's public scroll 620 or private scroll 622 if the participant decides to permit it. This decision may be implemented by means of a software flag, for example. This assures the participants that they can make annotations privately. It is

believed that this guarantee of privacy places the participants more at ease, and therefore more able to experiment or brainstorm. Certain alternative embodiments lack this feature. Certain other alternative embodiments can implement sessions either including or excluding this feature, at the option of the moderator. In these last embodiments, preferably some visual cue, such as an icon or a background color, is included in the participant's displays, so that they can readily identify whether the session permits such private annotations.

In certain embodiments, the moderator workstations and participant workstations are determined completely by software functionality. In these embodiments, more than one of the user's workstation can become the moderator workstation. For example, the moderator workstation could be selected by passing a security token from the moderator workstation to a participant workstation, thereby causing their roles to reverse. These embodiments could advantageously be used to facilitate the transfer of information between members of a group in environments where the transfer is not necessarily expected to flow primarily from one user to the others. For example, such embodiments could advantageously be used to facilitate information transfer in a remote conference of employees of one or more businesses. Note that it is not necessary in these embodiments for every workstation to have the same devices. For example, one workstation could be a preferred moderator workstation, having a plurality of display devices (for example, one for displaying the shared work area 330 and one for displaying the moderator private work area 340a), while the other workstations could be preferred participant workstations, having only a single display device each. Despite this limitation, the group work area could advantageously permit any of the participant work stations to become

the moderator work station. Such embodiments could advantageously be used, for example, in classroom settings, where the transfer of information is generally expected to predominantly flow from a designated moderator (the teacher), in order to facilitate situations in which a participant turns out to have a significant amount of information to convey to the other users. For example, in a class focused on architecture, the teacher might be especially interested in the perceptions of a foreign student about some aspect of foreign or domestic architecture. In that situation, the moderator could advantageously make the participant's workstation the moderator's workstation, so that the foreign student can augment his or her comments with images that are placed directly on each of the other user's public work areas.

While the invention has been illustrated and described in detail in the drawings and foregoing description, the same is to be considered as illustrative and not restrictive in character, it being understood that only the preferred embodiment has been shown and described and that all changes and modifications that come within the spirit of the invention are desired to be protected.